

Science on the Hill: Trinity ushers in new age of supercomputing

September 12, 2016

Science on the Hill: Trinity ushers in new age of supercomputing

by Gary Grider

As Los Alamos National Laboratory begins testing the second half of its new supercomputer, Trinity, the occasion highlights how intertwined scientific breakthroughs and computer innovations have become — and what a seminal and central role the Lab has played in that synergy.

Big Science, which today brings together theory, modeling, experiments that produce massive amounts of data and supercomputers to run incredibly sophisticated simulations providing feedback and validation to those theories and models, was largely pioneered at Los Alamos more than 70 years ago. When J. Robert Oppenheimer led his all-star team of scientists onto the Pajarito Plateau to unravel the secrets of the atom, they were embarking on an integrated research program at a scale the world had rarely seen.

Computers were there from the start. During World War II, the term “computers” applied to mathematicians — mostly women — who worked the differential equations by hand, with help from mechanical desktop calculators and simple punch-card machines from IBM. These were the first steps in the process of inventing how to use computers. Los Alamos scientists went on to run the first production job on the world’s first general-purpose electronic digital computer, ENIAC, and Nicholas Metropolis spearheaded development of the lab’s own computer, playfully dubbed MANIAC, in 1952, to continue the work of modeling nuclear processes.

Working with corporate collaborators, the lab has been stretching the boundaries of computing ever since, with innovation after innovation as the lab’s computers often topped the list of the fastest in the world. In 2008, the lab’s Roadrunner supercomputer became the first to break the petaflop barrier, processing a thousand million floating point operations each second. That kind of speed enables resolution sharpness in simulations that would have been unimaginable 70 years ago. In a global ocean climate model, for example, scientists can look at individual eddies in an ocean current.

None of these computers was “plug and play.” For each, the lab and its partners developed new software and hardware. Those innovations benefitted public and private computer users, offering solutions from how best to network very large clusters of computer processors to how to manage the data they produced.

Roadrunner's petaflop speed, for instance, was spinning out data at unprecedented rates during simulations running many months. Storage technology in that era struggled to keep up with the technology's ability to generate and consume data. During long-running calculations at very large scale, failures occur. To deal with recurring and somewhat random failure, an application periodically saves a snapshot of its current state. The program can restart from these checkpoints and continue for long periods.

If the stable storage that holds checkpoints is too slow, then computing time is lost either through spending too much time checkpointing, which bogs down the program, or by not checkpointing, which amplifies the effect of each failure.

The challenge intensified with Trinity, with its Cray architecture and two kinds of Intel processors. When fully installed, it will run about 40 times faster than Roadrunner and has memory roughly equal to the amount of memory of all the laptops in New Mexico. That performance would only make the checkpoint problems worse.

But several years ago, Los Alamos invented burst buffers, paving the way for Trinity. Using solid-state flash memory, similar to memory in the average smartphone, burst buffers take the rapid-fire data off the supercomputer processors and dole it out to slower disk drives while keeping the data handy for a restart. Performance improves, and flash memory for burst buffers is cheaper when bandwidth is taken into account compared to disk drives.

Burst buffers were installed for the first time on Trinity to support its crucial nuclear stockpile simulations. Other Department of Energy laboratories, academia, corporations and European supercomputer user sites are rapidly adopting this new technology. And the lab didn't stop with burst buffers — software engineers went on to develop another storage tool that allows supercomputers to save extremely large data sets for years on relatively inexpensive devices. Cloud-style inexpensive disk storage had not been applied to high-performance computing before.

Trinity and these storage tools continue the tradition of close collaboration between Los Alamos and computer vendors on the latest developments in computing technology. Big Science and its constant companion, Big Data, rely on the most advanced computers to simulate how the world works or to solve a mystery whose solution hides in a vast sea of data. The lab takes on challenges at a grand scale, from climate modeling to genetics, earthquakes to cancer, black holes to nuclear physics.

The computing innovations Los Alamos develops to solve these problems gives others the tools to address more everyday problems, such as by simulating a car crash as a means of improving real-world safety — research that ultimately enriches everyone's life.

Gary Grider is division leader of High Performance Computing at Los Alamos National Laboratory and a recognized international expert on supercomputing.

This article first appeared in the [Santa Fe New Mexican](#).